



GDM

DØ

Update on re-fixing and SAMv7 transition

28th Feb 2006

Gavin Davies, Imperial College London



Context

- Nov 05 - Reprocessing from raw of bulk of Run II data set with SAM-Grid complete
- Fixing - lighter weight correction, applied post (re-)processing
- Mid Dec 05 - Problem found with hadronic calorimeter correction

- Some of the relevant CY06 goals (as determined in fall 05)
 - ◆ Increased interoperability / automation \Rightarrow increased production rates along with increased functionality
 - ◆ Automated production of MC on LCG
 - ◆ Demonstration of primary processing with SAM-Grid
 - ◆ Deployment of SAM-Grid v7 (inc. d0runjob, MC, fixing etc)
 - ◆ SAM and SAM-Grid performance - inc data handling, deployment/operations - “deployment” team to start immediately
 - ◆ Automated production of MC with OSG




Re-fixing - I

- Data sets:
 - ◆ High priority Moriond skims - deadline of early / mid Jan
 - ◆ Full 1.4 billion events - from mid Jan - 1st March
 - ◆ Deadlines set by Physics needs
- Approach: Use all available techniques / resources to achieve above deadlines WHILST furthering long term goals where possible
- A great success - full 1.4 billion finished over 1 week early
 - ◆ First time fixing carried out using SAM-Grid
 - ◆ Increased interoperability
 - ◆ Use of CMS farm (OSG with local SAM-Grid installation)
 - ◆ Automated use of LCG facilities (without local SAM-Grid)
 - ◆ First use of OSG facilities without a local SAM-Grid installation
 - ◆ Improved data quality monitoring
- Thanks to many



Re-fixing - II

P17.09.03 Refix Status as of 24-Feb-2006 (all sites)



- CMS & WestGrid - “SAM-Grid” production on shared resources; LCG & OSG - “SAM-Grid” without local installation; CAB, farm, in2p3 - non-grid scripts.
- Flexibility allowed us to meet deadline and further long term goals
- Excellent “cross-disciplinary” effort made success possible
- Some ‘data-handling’ issues exposed (normal DH ops ~0.5FTE)
 - ◆ Overall through-put - access to servers, tape, off-site procedures...
 - ◆ Not a criticism - “expected” - part of the “operations/deployment” post
- Thanks
 - ◆ CMS: Lothar Bauderick, Ian Fisk, Jon Bakken
 - ◆ SAM-Grid team: Adam, Gabriele, Andrii, Parag, Robert
 - ◆ Other computing staff: Mike, Amber, Jason Allen, Joe Boyd..
 - ◆ DØ collaborators...



SAMv7 transition - I

- A big change - effects many areas -
 - ◆ Driven by both the experiment and the SAM team
- Dedicated project set-up in late Nov.
 - ◆ Areas affected and 1st pass time lines determined
 - ◆ Requires runjob (as opposed to “mc_runjob”)
 - ◆ Impacts SAMGrid (JIM) (60 calendar days (2 FTE) inc runjob integration)
 - ◆ Impacts on online system (few weeks)
 - ◆ Impacts on MC generation process (cardfiles, request system..)
 - ◆ Impacts on users - documentation, scripts, tools, accounting ...
- Aim: Address this, whilst maintaining other activities
 - ◆ Will be challenging - in particular maintaining operations.
“Operations/deployment” again ⇒ a dedicated post
 - ◆ Transition across sites will have to be adiabatic



SAMv7 transition - II

■ Status

- ◆ JIM changes largest (detailed workplan circulated to GDM before), basically on hold as uses same personnel (Andreii, Parag) as re-fixing - will now re-start.
- ◆ Online - Development basically complete, tests now needed
 - ◆ Used 'new' manpower ⇒ “deployment / operations” post
- ◆ Runjob development - brought in additional collaboration effort
 - ◆ Progressing (most executables in) but still considerable work
 - ◆ “Neglected” area - JIM / runjob / d0sandbox / d0rte
- ◆ MC - Cardfiles, request system, production page
 - ◆ Progress with tests on integration DB beginning
- ◆ Impact on users - scripts, tools, documentation...
 - ◆ Some progress, but again same set of core person(s)



Summary / Conclusions

- Full 1.4 billion events fixed ahead of schedule.
 - ◆ A tremendous success - thanks to all.
 - ◆ Advanced some of longer term goals, in particular on interoperability with LCG and OSG.

- An unavoidable consequence was delay to SAM v7 transition
 - ◆ Will now restart in earnest
 - ◆ Significant effort on both SAM-Grid & collaboration side needed
 - ◆ Timescales similar for both

- SAM v7 transition will leave “operations/deployment” exposed
 - ◆ Saw examples of this during re-fixing
 - ◆ New dedicated effort on this also covers limited data-handling support



Back-up slides



CY2006 Goals

Imperial College
London

(as defined in Fall 05 and previously presented to GDM)

- Use of luminosity database 1st Jan✓
- Automated production of MC on LCG Mid Jan
- Demonstration of primary processing with SAM-Grid End Jan
- Sustainable MC capacity of 5M events /week End Feb
- Completion of trigger database taskforce 1st Mar
- Deployment of SAM-Grid v7 (inc. d0runjob, MC and fixing) 1st Mar
- Entire Run IIa data (~1.3fb-1) set processed. 5th Mar
- p20 MC certified 19th Mar
- p20 MC in production on the farms 9th Apr
- CAB operating as part of Fermigrid / OSG 1st May
- SAM performance - increased cache metrics/ functionality 1st Jun
- SAM-Grid performance
 - ease of deployment (deployment team*), increased stability 1st Jun
- Primary processing with SAM-Grid as default Mid July
- Automated production of MC with OSG End Oct
- Sustainable MC capacity of 7M events /week 1st Nov
- Strategy for improved user analysis
 - grid user case?, resilient dcache? Mid Nov
- Automated submission of MC (or reprocessing) by SAM-shifter Mid Dec

* Activity also implicit in increased production rates - to start now.



- The following is the WBS of the project and a time estimate for its completion. Using the PERT Multi-point estimate rule, the total estimated FTE days for the project is (43 +- 3).
- 2 FTE can work in parallel during the development phase. Essentially no parallelization is achievable during the integration phase.
- Considering the 15-15 rule for working efficiency and week ends (holidays are not considered), the total estimated calendar days to complete the project is (60 +- 4). This estimate can extend to (72 +- 5) days if the porting of jim to python 2.4 fails and the contingency plan is enacted (see "contingency" section in the "Integrate system" item of the WBS).
- NOTES: Assumptions are noted next to the relevant activity in the WBS. All time estimates are expressed in FTE days and are rolled up using the PERT Multi-point estimate rule.



JIM Migration - II (Gabriele Garzoglio, Dec 02, 2005)

- - Port jim software suite to python 2.4.
- -- Perform show-stoppers tests (critical porting).
 - Time: (3.0 +- 0.7) (min 1, max 5, most-likely 3)
- --- xmldb_client (uses XML RCP libraries)
- --- xml_meta_configurator (uses pyxml intensively)
- -- Perform tests of the complete suite.
 - Time: (3.2 +- 0.5) (min 2, max 5, most-likely 3)
- - Implement required commands to sam v7
- -- Port "sam get dataset" command from sam v5.
 - Time: (2.0 +- 0.3) (min 1, max 3, most-likely 2)
- --- Lightweight consumer in poll mode
- --- Lightweight consumer in call-back mode (IS THIS NEEDED?)
- -- Add "poll"-mode to the station v6 routing interface (for LCG).
 - Time: (4.8 +- 0.8) (min 2, max 7, most-likely 5)
- - Port execution site software suite to sam v7.
- -- run_grid_job.py and plug in (jim_job_managers).
 - Time: (2.0 +- 0.3) (min 1, max 3, most-likely 2)
- -- wrappers (jim_job_managers): check the use of the sam command line commands.
 - Time: (1.2 +- 0.2) (min 1, max 2, most-likely 1)
- - Port cliet site software suite to sam v7.
 - Time: (3.0 +- 0.3) (min 2, max 4, most-likely 3)
- -- Port "check job consistency" modules (interaction with db server)



JIM Migration - III (Gabriele Garzoglio, Dec 02, 2005)

- -Integrate system
- -- Integrate JIM (new sam_client software distribution).
- Time: (1.2 +- 0.2) (min 1, max 2, most-likely 1)
- -- [Contingency: What if we cannot port JIM to python v2.4?
- Distribute different python versions for JIM and Runjob to the worker node.
- Time: (min 4, max 10, most-likely 6)
-]
- -- Integrate Runjob.
- Time: (14.5 +- 2.2) (min 7, max 20, most-likely 15)
- (ASSUMES:
- o Runjob integrated with SAM for data storing and MC request system
- o Runjob developers presence
-)
- -- Integration tests.
- Time: (8.5 +- 0.8) (min 5, max 10, most-likely 9)



Deployment Team - Draft Charge

SAMGrid is being increasingly used by Tevatron experiments, and in particular by the DØ collaboration. DØ is effectively reliant on SAMGrid for much of its production computing. SAM has been used for all data handling since the start of RunII, SAMGrid is now the default for Monte Carlo production and recent developments include the reprocessing and fixing of data. Next steps include the inclusion of other production tasks such as primary processing or skimming and finally individual analysis jobs. In parallel there is a very active programme to make SAMGrid fully interoperable with other grids such as LCG and OSG, to give the Run II experiments access to the resources required to analysis the growing Run II data sets.

In the past the same developers have been responsible for assisting with the installation at remote sites, operational support and optimization of the overall system performance. In view of the increased functionality, increased scale and ongoing development this is no longer tenable. Hence a SAMGrid deployment / operations team has been formed. This group, as a sub-team within the SAMGrid project, will be responsible for:

- ◆ Determining and documenting the appropriate deployment procedure for SAMGrid, either as a native installation or as part of the LCG and OSG interoperability projects.
- ◆ Assisting with the deployment of SAMGrid at remote sites, either as a native installation or as part of the LCG and OSG interoperability projects.
- ◆ Maintenance of the central SAMGrid infrastructure at Fermilab at the appropriate production level.
- ◆ Assisting with the operational support of the remote facilities in conjunction with the remote sites themselves.
- ◆ Monitoring and optimizing the overall data-handling associated with SAMGrid operations, in conjunction with the experiment's existing data-handling support and the remote sites. This will involve determining whether suitable tools are available and contributing to their development if not.
- ◆ The precise nature of these tasks will evolve as the team is established. The team will have a SAM-Grid core, supplemented with a member from each collaboration on a 6 month rotation at Fermilab.